

Adaptive Interventions with User-Defined Goals for Health Behavior Change

Aishwarya Mandyam, Matthew Jörke, Barbara E. Engelhardt, Emma Brunskill

March 28th, 2024

1 Introduction

Physical inactivity is a global health crisis linked to severe health issues like cardiovascular disease and type-2 diabetes. With over a quarter of the world’s population not meeting recommended physical activity levels, there’s a pressing need for effective physical activity interventions. Mobile health (mHealth) applications offer a promising solution due to their accessibility and scalability. However, these interventions often fall short in effect size and user adherence when compared to traditional human coaching. One missing piece in these digital interventions is personalized goal-setting, a proven strategy in health coaching for building sustainable physical activity habits. Our research integrates individualized goal-setting into adaptive algorithms for mHealth interventions, aiming to enhance their effectiveness across a large population of users that have different goals and lifestyle constraints.

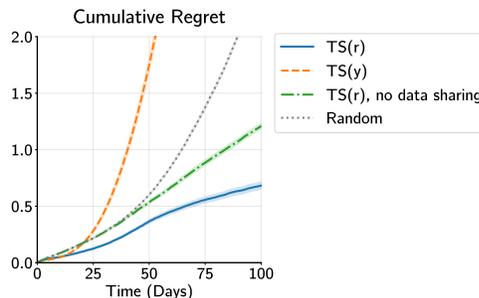
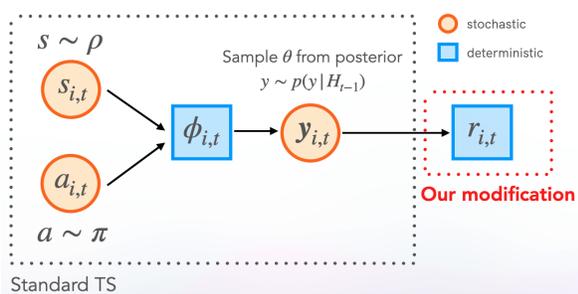
2 Methods

We propose a modification to the Thompson sampling algorithm, which is commonly used to learn adaptive intervention policies, to optimize for personalized reward functions that reflect individual goals. Instead of optimizing for outcomes, such as step count, our algorithm optimizes for a personalized reward function that is modeled as a weighted sum of utility functions. Each utility function measures the change in a user’s personal satisfaction based on their daily physical activity and factors like the number of notifications they receive. This approach not only acknowledges the unique preferences of each user but also leverages the shared structure across individuals to enhance learning efficiency. We developed a physical activity simulator to test our algorithm, incorporating factors like treatment effect heterogeneity and notification burden to closely mimic real-world scenarios.

3 Results

We demonstrate results in a step count simulator, which is modeled based on the HeartSteps and the MyHeart Counts studies. Our modified Thompson sampling algorithm outperformed traditional methods that either don’t share data across users or fail to consider individualized rewards. Specifically, our approach achieved lower cumulative regret, meaning it was more effective in meeting users’ personalized goals. This balance is crucial for maintaining engagement and promoting

long-term health behavior change. Furthermore, our findings suggest that sharing data across users allows us to more quickly learn optimal intervention policies for each individual’s unique preferences and goals.



(a) A visual diagram of our modification to the Thompson Sampling algorithm.

(b) Our approach, TS(r), achieves the lowest cumulative regret in comparison to the baselines.

4 Conclusion

Our research represents an advancement in the design of adaptive mHealth interventions. By incorporating user-defined goals into the algorithmic framework, we’ve demonstrated a path to more effective and engaging digital health solutions. This approach not only respects the individuality of users but also maintains the scalability and cost-effectiveness of mHealth applications. Future work will focus on refining our model with semi-synthetic simulators based on real-world data and reducing restrictions on the parameterization of personalized reward functions. Ultimately, our goal is to bridge the gap between the personalized touch of human coaching and the scalability of digital interventions, making health behavior change accessible and effective for everyone.